Terms used **normalization sort partial sum** Found **43** of **198,146**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 43 Result page: **1** 2 3 next

Relevance scale

1 Approximate computation of multidimensional aggregates of sparse data using wavelets
Jeffrey Scott Vitter, Min Wang
June 1999 **ACM SIGMOD Record , Proceedings of the 1999 ACM SIGMOD international conference on Management of data SIGMOD '99**, Volume 28 Issue 2
Publisher: ACM Press
Full text available: pdf(1.67 MB) Additional Information: full citation, abstract, references, citings, index terms

Computing multidimensional aggregates in high dimensions is a performance bottleneck for many OLAP applications. Obtaining the exact answer to an aggregation query can be prohibitively expensive in terms of time and/or storage space in a data warehouse environment. It is advantageous to have fast, approximate answers to OLAP aggregation queries. In this paper, we present a novel method that provides approximate answers to high-dimensional OLAP aggregation queries in massive spars ...

2 Algorithmic transforms for efficient energy scalable computation
Amit Sinha, Alice Wang, Anantha P. Chandrakasan
August 2000 **Proceedings of the 2000 international symposium on Low power electronics and design ISLPED '00**
Publisher: ACM Press
Full text available: pdf(171.16 KB) Additional Information: full citation, abstract, references, citings, index terms

We introduce the notion of energy scalable computation on general purpose processors. The principle idea is to maximize computational qualityfor a given energy constraint. Teh desirable energy-quality behavior of algorithms is discussed. subsequently the energy-quality scalability of three distinct categories of commonly used signal processing algorithms (viz. filtering, frequency domain transforms and classification) are analyzed on the StrongARM SA-1100 processor and transformations are d ...

3 Selectivity estimators for multidimensional range queries over real attributes
Dimitrios Gunopulos, George Kollios, J. Tsotras, Carlotta Domeniconi
April 2005 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 14 Issue 2
Publisher: Springer-Verlag New York, Inc.
Full text available: pdf(321.96 KB) Additional Information: full citation, abstract

Estimating the selectivity of multidimensional range queries over real valued attributes

has significant applications in data exploration and database query optimization. In this paper, we consider the following problem: given a table of *d* attributes whose domain is the real numbers and a query that specifies a range in each dimension, find a good approximation of the number of records in the table that satisfy the query. The simplest approach to tackle this problem is to assume that the ...

4 Optimizing computations in a transposed direct form realization of floating-point LTI-FIR systems

N. Sankarayya, K. Roy, D. Bhattacharya

November 1997 **Proceedings of the 1997 IEEE/ACM international conference on Computer-aided design ICCAD '97**

Publisher: IEEE Computer Society

Full text available: pdf(188.36 KB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

The inherent computational redundancy in discrete-time LTI-FIR system response computations in Digital Signal Processing have been exploited in a variety of ways to minimize the computational complexity. We present an improved algorithm-level computational optimization that uses sorted recursive differences between coefficients representing the system transfer function with a Floating-Point number representation to extract maximum benefits from this redundancy. It can be applied to any LTI-FIR s ...

Keywords: DSP, FIR, filters, low-power, Differential Coefficients, Sorted Recursive Differences, multiplication

5 On accurate floating-point summation

Michael A. Malcolm

November 1971 **Communications of the ACM**, Volume 14 Issue 11

Publisher: ACM Press

Full text available: pdf(478.31 KB) Additional Information: full citation, abstract, references, citings

cumulation of floating-point sums is considered on a computer which performs t-digit base &bgr; floating-point addition with exponents in the range —m to M. An algorithm is given for accurately summing n t-digit floating-point numbers. Each of these n numbers is split into q parts, forming q·n t-digit floating-point numbers. Each of ...

Keywords: error analysis, floating-point summation

6 Remark on algorithm 347: An efficient algorithm for sorting with minimal storage

Richard Peto

October 1970 **Communications of the ACM**, Volume 13 Issue 10

Publisher: ACM Press

Full text available: pdf(748.18 KB) Additional Information: full citation, citings

Keywords: digital computer sorting, minimal storage sorting, ranking, sorting

7 When indexing equals compression: Experiments with compressing suffix arrays and applications

Luca Foschini, Roberto Grossi, Ankur Gupta, Jeffrey Scott Vitter

October 2006 **ACM Transactions on Algorithms (TALG)**, Volume 2 Issue 4

Publisher: ACM Press
Full text available: pdf(351.20 KB) Additional Information: full citation, abstract, references, index terms

We report on a new experimental analysis of high-order entropy-compressed suffix arrays, which retains the theoretical performance of previous work and represents an improvement in practice. Our experiments indicate that the resulting text index offers state-of-the-art compression. In particular, we require roughly 20% of the original text size---without requiring a separate instance of the text. We can additionally use a simple notion to encode and decode block-sorting transforms (such a ...

Keywords: Burrows--Wheeler Transform, Entropy, suffix array, text indexing

8 P2P and network algorithms: Efficient top-K query calculation in distributed networks
Pei Cao, Zhe Wang
July 2004 **Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing PODC '04**
Publisher: ACM Press
Full text available: pdf(206.90 KB) Additional Information: full citation, abstract, references, citings, index terms

This paper presents a new algorithm to answer top-*k* queries (e.g. "find the *k* objects with the highest aggregate values") in a distributed network. Existing algorithms such as the Threshold Algorithm [10] consume an excessive amount of bandwidth when the number of nodes, *m*, is high. We propose a new algorithm called "Three-Phase Uniform Threshold" (TPUT). TPUT reduces network bandwidth consumption by pruning away ineligible objects, and terminates in three round-trips regard ...

Keywords: distributed networks, instance optimality, top-k algorithms

9 Efficiently supporting ad hoc queries in large datasets of time sequences
Flip Korn, H. V. Jagadish, Christos Faloutsos
June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data SIGMOD '97**, Volume 26 Issue 2
Publisher: ACM Press
Full text available: pdf(1.43 MB) Additional Information: full citation, abstract, references, citings, index terms

Ad hoc querying is difficult on very large datasets, since it is usually not possible to have the entire dataset on disk. While compression can be used to decrease the size of the dataset, compressed data is notoriously difficult to index or access. In this paper we consider a very large dataset comprising multiple distinct time sequences. Each point in the sequence is a numerical value. We show how to compress such a dataset into a format that supports ad hoc querying, provided ...

10 Parallel processing image synthesis and anti-aliasing
Richard Weinberg
August 1981 **ACM SIGGRAPH Computer Graphics , Proceedings of the 8th annual conference on Computer graphics and interactive techniques SIGGRAPH '81**, Volume 15 Issue 3
Publisher: ACM Press
Full text available: pdf(834.55 KB) Additional Information: full citation, abstract, references, citings, index terms

The continuing evolution of microelectronics provides the tools for developing new methods of synthesizing digital images by utilizing parallel processing architectures which hold the promise of reliability, flexibility and low cost. Beginning with the earliest real-

time flight simulators, parallel processing architectures for image synthesis have been built, but "anti-aliasing" remains a problem. A parallel processing architecture is described and simulated which consists of a serial chain ...

Keywords: Aliasing, Computer graphics, Parallel processing

11 Pipeline Architecture

C. V. Ramamoorthy, H. F. Li
March 1977 **ACM Computing Surveys (CSUR)**, Volume 9 Issue 1
Publisher: ACM Press
Full text available: pdf(3.53 MB) Additional Information: full citation, references, citings, index terms

12 A defect tolerant self-organizing nanoscale SIMD architecture

Jaidev P. Patwardhan, Vijeta Johri, Chris Dwyer, Alvin R. Lebeck
October 2006 **ACM SIGOPS Operating Systems Review , ACM SIGPLAN Notices , ACM SIGARCH Computer Architecture News , Proceedings of the 12th international conference on Architectural support for programming languages and operating systems ASPLOS-XII**, Volume 40 , 41 , 34 Issue 5 , 11 , 5
Publisher: ACM Press
Full text available: pdf(633.08 KB) Additional Information: full citation, abstract, references, index terms

The continual decrease in transistor size (through either scaled CMOS or emerging nano-technologies) promises to usher in an era of tera to peta-scale integration. However, this decrease in size is also likely to increase defect densities, contributing to the exponentially increasing cost of top-down lithography. Bottom-up manufacturing techniques, like self assembly, may provide a viable lower-cost alternative to top-down lithography, but may also be prone to higher defects. Therefore, regardle ...

Keywords: DNA, SIMD, bit-serial, data parallel, defect tolerance, nanocomputing, self-organizing

13 On the type structure of standard ML

Robert Harper, John C. Mitchell
April 1993 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 15 Issue 2
Publisher: ACM Press
Full text available: pdf(2.68 MB) Additional Information: full citation, references, citings, index terms, review

14 Session 1C: methods for DSP synthesis and debugging: FIR filter synthesis algorithms for minimizing the delay and the number of adders

Hyeong Ju Kang, Hansoo Kim, In Cheol Park
November 2000 **Proceedings of the 2000 IEEE/ACM international conference on Computer-aided design ICCAD '00**
Publisher: IEEE Press
Full text available: pdf(65.12 KB) Additional Information: full citation, abstract, references

As the complexity of digital filters is dominated by the number of multiplications, many works have focused on minimizing the complexity of multiplier blocks that compute the constant coefficient multiplications required in filters. Although the complexity of

multiplier blocks is significantly reduced by using efficient techniques such as decomposing multiplications into simple operations and sharing common subexpressions, previous works have not considered the delay of multiplier blocks which i ...

15 Remark on algorithm 368: Numerical inversion of Laplace transforms

Harald Stehfest
October 1970 **Communications of the ACM**, Volume 13 Issue 10
Publisher: ACM Press
Full text available: pdf(748.18 KB) Additional Information: full citation

Keywords: Laplace transform inversion, integral equations, integral transformations

16 Remark on algorithm 304: Normal curve integral

Bo Holmgren
October 1970 **Communications of the ACM**, Volume 13 Issue 10
Publisher: ACM Press
Full text available: pdf(748.18 KB) Additional Information: full citation

Keywords: normal curve integral, probability, special functions

17 Algorithm 400: Modified Havie integration

George C. Wallick
October 1970 **Communications of the ACM**, Volume 13 Issue 10
Publisher: ACM Press
Full text available: pdf(748.18 KB) Additional Information: full citation, references, citings

Keywords: Havie integration, Romberg quadrature, modified Romberg-quadrature, numerical integration, rectangle values, trapezoid values

18 Algorithm 399: Spanning tree

Jouko J. Seppänen
October 1970 **Communications of the ACM**, Volume 13 Issue 10
Publisher: ACM Press
Full text available: pdf(748.18 KB) Additional Information: full citation, references

Keywords: graph, spanning tree, tree

19 Algorithm 398: Tableless date conversion

Richard A. Stone
October 1970 **Communications of the ACM**, Volume 13 Issue 10
Publisher: ACM Press
Full text available: pdf(748.17 KB) Additional Information: full citation

Keywords: calendar, date

20 Algorithm 397: An integer programming problem

S. K. Chang, A. Gill
October 1970 **Communications of the ACM**, Volume 13 Issue 10
Publisher: ACM Press
Full text available: pdf(748.17 KB) Additional Information: full citation, references

Keywords: change-making problem, integer programming

Results 1 - 20 of 43 Result page: **1** 2 3 next

Terms used **normalization sort partial sum** Found **43** of **198,146**

Sort results by relevance
Display results expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 21 - 40 of 43 Result page: previous 1 **2** 3 next

Relevance scale

21 Algorithm 396: Student's t-quantiles
G. W. Hill
October 1970 **Communications of the ACM**, Volume 13 Issue 10
Publisher: ACM Press
Full text available: pdf(748.17 KB) Additional Information: full citation, references

Keywords: Student's t-statistic, asymptotic approximation, quantile

22 Algorithm 395: Students t-distribution
G. W. Hill
October 1970 **Communications of the ACM**, Volume 13 Issue 10
Publisher: ACM Press
Full text available: pdf(748.17 KB) Additional Information: full citation, references

Keywords: approximation, asymptotic expansion, distribution function, student's t-statistic

23 ACM Algorithm 396: Student's t-Quantiles
G. W. Hill
October 1970 **Communications of the ACM**, Volume 13 Issue 10
Publisher: ACM Press
Full text available: pdf(748.17 KB) Additional Information: full citation, references

Keywords: asymptotic expansion, quantile, student's t-statistic

24 ACM Algorithm 395: Student's t-distribution
G. W. Hill
October 1970 **Communications of the ACM**, Volume 13 Issue 10

Publisher: ACM Press
Full text available: pdf(748.17 KB) Additional Information: full citation, references, citings

Keywords: approximation, asymptotic expansion, distribution function, students t-statistic

25 The efficiency of inverted index and cluster searches

Ellen M. Voorhees
September 1986 **Proceedings of the 9th annual international ACM SIGIR conference on Research and development in information retrieval SIGIR '86**
Publisher: ACM Press
Full text available: pdf(891.67 KB) Additional Information: full citation, abstract, references, citings

The processing time and disk space requirements of an inverted index and top-down cluster search are compared. The cluster search is shown to use both more time and more disk space, mostly due to the large number of cluster centroids needed by the search. When shorter centroids are used, the efficiency of the cluster search improves, but the inverted index search remains more efficient.

26 Validation of Scientific Programs

William E. Howden
June 1982 **ACM Computing Surveys (CSUR)**, Volume 14 Issue 2
Publisher: ACM Press
Full text available: pdf(2.92 MB) Additional Information: full citation, references, citings, index terms

27 An efficient boosting algorithm for combining preferences

Yoav Freund, Raj Iyer, Robert E. Schapire, Yoram Singer
December 2003 **The Journal of Machine Learning Research**, Volume 4
Publisher: MIT Press
Full text available: pdf(392.20 KB) Additional Information: full citation, abstract, citings, index terms

We study the problem of learning to accurately rank a set of objects by combining a given collection of ranking or preference functions. This problem of combining preferences arises in several applications, such as that of combining the results of different search engines, or the "collaborative-filtering" problem of ranking movies for a user based on the movie rankings provided by other users. In this work, we begin by presenting a formal framework for this general problem. We then describe and ...

28 The height of a random binary search tree

Bruce Reed
May 2003 **Journal of the ACM (JACM)**, Volume 50 Issue 3
Publisher: ACM Press
Full text available: pdf(232.76 KB) Additional Information: full citation, abstract, references, citings, index terms

Let Hn be the height of a random binary search tree on n nodes. We show that there exist constants $\alpha = 4.311...$ and $\beta = 1.953...$ such that $\mathbf{E}(Hn) = \alpha \ln n - \beta \ln \ln n + O(1)$, We also show that $\mathbf{Var}(Hn) = O(1)$.

Keywords: Binary search tree, asymptotics, height, probabilistic analysis, random tree, second moment method

29 Research sessions: query processing II: Efficient k-NN search on vertically decomposed data

Arjen P. de Vries, Nikos Mamoulis, Niels Nes, Martin Kersten

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02**

Publisher: ACM Press

Full text available: pdf(1.26 MB) Additional Information: full citation, abstract, references, citings, index terms

Applications like multimedia retrieval require efficient support for similarity search on large data collections. Yet, nearest neighbor search is a difficult problem in high dimensional spaces, rendering efficient applications hard to realize: index structures degrade rapidly with increasing dimensionality, while sequential search is not an attractive solution for repositories with millions of objects. This paper approaches the problem from a different angle. A solution is sought in an unconvent ...

30 Algorithm 719; Multiprecision translation and execution of FORTRAN programs

David H. Bailey

September 1993 **ACM Transactions on Mathematical Software (TOMS)**, Volume 19 Issue 3

Publisher: ACM Press

Full text available: pdf(2.03 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper describes two Fortran utilities for multiprecision computation. The first is a package of Fortran subroutines that perform a variety of arithmetic operations and transcendental functions on floating point numbers of arbitrarily high precision. This package is in some cases over 200 times faster than that of certain other packages that have been developed for this purpose. The second utility is a translator program, which facilitates the conversion of ordinary Fortran p ...

Keywords: multiple-precision computation, multiprecision arithmetic

31 How tall is a tree?

Bruce Reed

May 2000 **Proceedings of the thirty-second annual ACM symposium on Theory of computing STOC '00**

Publisher: ACM Press

Full text available: pdf(496.07 KB) Additional Information: full citation, references, citings, index terms

32 Interactive multiresolution surface viewing

Andrew Certain, Jovan Popovic, Tony DeRose, Tom Duchamp, David Salesin, Werner Stuetzle

August 1996 **Proceedings of the 23rd annual conference on Computer graphics and interactive techniques SIGGRAPH '96**

Publisher: ACM Press

Full text available: pdf(420.60 KB) Additional Information: full citation, references, citings, index terms

Keywords: geometric modeling, multiresolution analysis, texture mapping, viewer, wavelets

33 Session 8A: When indexing equals compression: experiments with compressing

suffix arrays and applications
Roberto Grossi, Ankur Gupta, Jeffrey Scott Vitter
January 2004 **Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms SODA '04**
Publisher: Society for Industrial and Applied Mathematics
Full text available: pdf(229.91 KB) Additional Information: full citation, abstract, references, citings

We report on a new and improved version of high-order entropy-compressed suffix arrays, which has theoretical performance guarantees similar to those in our earlier work [16], yet represents an improvement in practice. Our experiments indicate that the resulting text index offers state-of-the-art compression. In particular, we require roughly 20% of the original text size---without requiring a separate instance of the text---and support fast and powerful searches. To our knowledge, this is the b ...

34 Symbolic algebra in theoretical seismology
E. Nyland, C. H. Chapman
March 1971 **Proceedings of the second ACM symposium on Symbolic and algebraic manipulation SYMSAC '71**
Publisher: ACM Press
Full text available: pdf(1.17 MB) Additional Information: full citation, abstract, references, index terms

The solution of those systems of partial differential equations which arise in elastic wave propagation can be aided in some respects by the use of symbolic algebra manipulations on a digital computer. A serious difficulty in the algebra is the lack of ability torecognize more than trivial simplifications. It is possible to manipulate Legendre and Bessel functions with confidence, but algebraic swell leads to exorbitant core requirements. Investigation of algebraic details of aspects of a p ...

35 Comparing Bayes model averaging and stacking when model approximation error cannot be ignored
Bertrand Clarke
December 2003 **The Journal of Machine Learning Research**, Volume 4
Publisher: MIT Press
Full text available: pdf(248.77 KB) Additional Information: full citation, abstract, references, index terms

We compare Bayes Model Averaging, BMA, to a non-Bayes form of model averaging called stacking. In stacking, the weights are no longer posterior probabilities of models; they are obtained by a technique based on cross-validation. When the correct data generating model (DGM) is on the list of models under consideration BMA is never worse than stacking and often is demonstrably better, provided that the noise level is of order commensurate with the coefficients and explanatory variables. Here, howe ...

36 The contribution to performance of instruction set usage in System/370
O. R. LaMaire, W. W. White
November 1986 **Proceedings of 1986 ACM Fall joint computer conference ACM '86**
Publisher: IEEE Computer Society Press
Full text available: pdf(1.18 MB) Additional Information: full citation, references, citings, index terms

37 Relevant context inference
Ramkrishna Chatterjee, Barbara G. Ryder, William A. Landi
January 1999 **Proceedings of the 26th ACM SIGPLAN-SIGACT symposium on Principles of programming languages POPL '99**
Publisher: ACM Press
Full text available: pdf(1.93 MB) Additional Information: full citation, references, citings, index terms

38 An analysis of the Cray-1 computer

Richard L. Sites
April 1978 **Proceedings of the 5th annual symposium on Computer architecture ISCA '78**
Publisher: ACM Press

Full text available: pdf(687.82 KB) Additional Information: full citation, abstract, references, citings, index terms

The Cray-1 is an extremely high-speed computer, intended to be used for large floating-point scientific computations. However, it is a well-balanced machine that can gracefully be used on a wide class of problems. The machine has two major architectural innovations: (1) 128 backup registers which represent a new layer in the memory hierarchy, essentially a programmer or compiler-managed cache, and (2) 8 vector registers holding up to 64 words each, and operated on by vector instructions. In ...

39 Scalar fused multiply-add instructions produce floating-point matrix arithmetic provably accurate to the penultimate digit

Yves Nievergelt
March 2003 **ACM Transactions on Mathematical Software (TOMS)**, Volume 29 Issue 1
Publisher: ACM Press

Full text available: pdf(219.83 KB) Additional Information: full citation, abstract, references, citings, index terms

Combined with doubly compensated summation, scalar fused multiply-add instructions redefine the concept of floating-point arithmetic, because they allow for the computation of sums of real or complex matrix products accurate to the penultimate digit. Particular cases include complex arithmetic, dot products, cross products, residuals of linear systems, determinants of small matrices, discriminants of quadratic, cubic, or quartic equations, and polynomials.

Keywords: Doubly compensated summation, floating-point arithmetic, fused multiply-add instruction, matrix arithmetic, provable accuracy, rounding error

40 Passage retrieval revisited

Marcin Kaszkiel, Justin Zobel
July 1997 **ACM SIGIR Forum , Proceedings of the 20th annual international ACM SIGIR conference on Research and development in information retrieval SIGIR '97**, Volume 31 Issue SI
Publisher: ACM Press

Full text available: pdf(1.35 MB) Additional Information: full citation, abstract, references, citings, index terms

Ranking based on passages addresses some of the shortcomings of whole-document ranking. It provides convenient units of text to return to the user, avoids the difficulties of comparing documents of different length, and enables identification of short blocks of relevant material amongst otherwise irrelevant text. In this paper we explore the potential of passage retrieval, based on an experimental evaluation of the ability of passages to identify relevant documents. We compare our sch ...

Terms used **normalization sort partial sum** Found **43** of **198,146**

Sort results by relevance

Display results expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 41 - 43 of 43 Result page: previous 1 2 **3**

Relevance scale

41 Supercomputing and transputers

Falk Langhammer, Francis Wray

August 1992 **Proceedings of the 6th international conference on Supercomputing ICS '92**

Publisher: ACM Press

Full text available: pdf(1.72 MB) Additional Information: full citation, abstract, references, index terms

It will be studied which degree parallel supercomputers can be scaled to. Necessary measures to achieve a maximum scalability will be discussed, and a case-study be presented. To this purpose, a new class of "supermassively parallel architectures" is introduced, and the notation of scalable architectures will be extended to reflect the impact of technological progress onto cost-functions. For systems in this class, the performance efficiency of applications is discussed and two ...

42 Simulating (logcn)-wise independence in NC

Bonnie Berger

October 1991 **Journal of the ACM (JACM)**, Volume 38 Issue 4

Publisher: ACM Press

Full text available: pdf(1.28 MB) Additional Information: full citation, references, citings, index terms

Keywords: discrepancy, removing randomness

43 Bandwidth and traffic estimation techniques: Single-hop probing asymptotics in available bandwidth estimation: sample-path analysis

Xiliang Liu, Kaliappa Ravindran, Benyuan Liu, Dmitri Loguinov

October 2004 **Proceedings of the 4th ACM SIGCOMM conference on Internet measurement IMC '04**

Publisher: ACM Press

Full text available: pdf(420.64 KB) Additional Information: full citation, abstract, references, citings, index terms

In this paper, we take the sample-path approach in analyzing the asymptotic behavior of single-hop bandwidth estimation under bursty cross-traffic and show that these results are provably different from those observed under fluid models of prior work. This difference, which we call the probing bias, is one of the previously unknown factors that can cause measurement inaccuracies in available bandwidth estimation. We present an

analytical formulation of "packet probing," based on which we deri ...

Keywords: bandwidth measurement, packet train probing

Results 41 - 43 of 43 Result page: previous 1 2 **3**